From automation to autonomy

A pragmatic framework for the AI-native service provider

Telecommunications service providers must build an autonomous intelligent network that can grow at the speed of innovation to meet growing customer demand.

Operational complexity: The roadblock to new services

Modern telecommunication networks have reached a level of complexity that renders manual, script-based operations obsolete. As telecommunication service providers adopt 5G standalone network core architectures and multivendor radio access networks (RANs), they are managing billions of daily events across hundreds of thousands of network elements.

In today's telecommunication industry, automation is no longer a competitive advantage, but a foundational requirement. Many telecommunication service providers now accept the need for a qualitative shift toward self-healing, intent-driven systems. The business case is compelling: Telecommunication service providers see that speeding time to market is the primary motivator for autonomy, while early adopters acknowledge a significant reduction in operating expenses and faster mean time to repair (MTTR).

Transitioning to an open source-based, intelligent, and autonomous network addresses 3 critical executive priorities:

- **Creating new revenue:** By moving away from rigid, proprietary, and isolated structures, telecommunication service providers can launch Network-as-a-Service (NaaS) and dynamic 5G slicing. With an open architecture, they can integrate third-party AI applications and edge services, turning the network into a revenue-generating platform.
- **Reducing costs:** Telecommunication service providers who choose open source can opt for cost-effective, standardized hardware. Additionally, automated energy management and intelligent workload placement can significantly lower power usage.
- **Mitigating risk:** Open source models make complying with strict regulations easier because they are inherently transparent and explainable. Telecommunication service providers maintain digital sovereignty by keeping training data and model weights on-premise, ensuring their competitive intelligence remains private.

Levels of autonomy: Charting your path with the TM Forum model

The journey to full autonomy is governed by TM Forum's 6-level maturity model (levels 0–5). Most telecommunication service providers currently operate at level 2 (partial automation), with the goal of reaching level 4 (high automation), where the network detects environmental changes and corrects itself without human intervention.

7 principles of the TM Forum maturity model

TM Forum's maturity model relies on:

- **Decoupled operational layers:** Runs business, service, and resource layers independently.
- **Intent-driven interfaces:** Communicates what to achieve, not how to achieve it.
- **Closed-loop automation:** Includes continuous cycles of awareness, analysis, decision, and execution (AADE).
- **Endogenous intelligence:** Embeds AI within network elements, not just in a central cloud.
- **Single-domain autonomy:** Divides large networks into self-governing units (e.g., RAN or core).
- **Cross-domain collaboration:** Institutes horizontal federation across domains.
- **Multilevel interworking:** Operates consistently across different maturity levels.

Most telecommunication service providers start using partial automation (level 2), then move to conditional automation (level 3), with their sights set on high automation (level 4) and full automation (level 5).

3-layer architecture

TM Forum's architecture follows 3 foundational principles:

- **Business operations:** Manages customer lifecycle, product offerings, and business-level intents.
- **Service operations:** Turns business objectives into action items, coordinates across domains, and manages end-to-end service-level agreement (SLA) assurance.
- **Resource operations:** Includes technology domains—such as RAN, core, transport, and edge—that run their own AADE closed loops. Domain controllers manage cloud-native network functions (CNFs) by translating service intents into network configurations.

The knowledge and intelligence plane supports these 3 layers, providing the AI brain that feeds trained models, analytics, and application catalogs to all 3 operational layers.

Activating these layers is the main goal. TM Forum defines this goal through the intent common model (TR290), which provides machine-readable specifications for what telecommunication service providers want their networks to achieve.

The open foundation for AI-powered autonomy

Red Hat provides a reliable open source execution engine that works with the service operations and resource operations layers of the TM Forum's model, as shown in Table 1.

Table 1: Red Hat solutions mapping to the TM Forum architecture

TM Forum layer	Red Hat solution	Strategic value
Business operations	Red Hat® OpenShift® (as a platform)	Works with selected ecosystem partners
Service operations	Red Hat Advanced Cluster Management for Kubernetes, Red Hat AI (Red Hat Enterprise Linux® AI, Red Hat OpenShift AI, and Red Hat Inference Server)	Manages cross-domain orchestration and provides the AI brain for SLA assurance
Resource operations	Red Hat OpenShift, Red Hat OpenShift Virtualization, Red Hat Enterprise Linux, Red Hat Ansible® Automation Platform, Event-Driven Ansible, and automation intelligence assistant	Reliable execution of CNFs and localized site intelligence at the edge for real-time response

Red Hat's portfolio uses a consistent operational model that replaces disparate, proprietary systems:

- **The execution engine (Red Hat Ansible Automation Platform)** implements configuration changes with automated playbooks. Event-Driven Ansible closes the loop by triggering remediations in minutes, not hours.
- **The AI brain (Red Hat AI)** implements the knowledge and intelligence plane. With Red Hat AI, telecommunication service providers can fine-tune open foundation models on proprietary network data without compromising data sovereignty.
- **The action layer (Red Hat OpenShift)** acts as the common cloud infrastructure for autonomous domains (RAN, core, and edge) and supports site intelligence by serving lightweight models at the edge of the network.

Business and service intents must be expressed, separated, and routed to the right domains, with Red Hat Ansible Automation Platform providing the automation execution engine that brings intent to life.

Faster, smarter, and leaner: The measurable impact of intelligent autonomy

The following real-world use cases are being solved with intelligent autonomous networks that map to AADE.

Elevated latency affecting enterprise customers

A 5G core network function begins experiencing high levels of latency:

- **Awareness:** Monitoring software like Prometheus (running on Red Hat OpenShift) detects the anomaly via metrics from the 5G core CNFs. Event-Driven Ansible receives the alert.
- **Analysis:** The alert triggers an AI agent (built using Red Hat AI) that queries the knowledge graph for topology context, pulls recent change logs, and runs root cause analysis (RCA) using an open source, fine-tuned model served by Red Hat Inference Server.
- **Decision:** The AI agent identifies the root cause—a 5G core user plane function (UPF) scaling policy that was misconfigured after a recent update. The AI agent generates a remediation recommendation and a confidence score. Once the confidence exceeds the policy threshold of more than 90%, it proceeds without human intervention.
- **Execution:** Event-Driven Ansible executes the remediation playbook. It rolls back the UPF scaling policy, verifies the latency returns to SLA compliance, and generates an incident report.

Red Hat AI is a platform optimized for production AI workloads that removes the integration burden. Instead of assembling individual components, telecommunication service providers get a reliable AI platform that is ready for autonomous intelligent network demands.

Level-3 customer service assistant for network operations

This provides expert-level network troubleshooting and technical support. It bridges the gap between level-1 and level-2 agents and level-3 engineers:

- **Awareness:** A ticket highlighting customer experience issues goes through lower-level support teams, which triggers the agentic workflow built using Red Hat AI. Typically, it would be triggered manually within an IT service management ticketing system.
- **Analysis:** The alert triggers an AI leader orchestrator agent (built using Red Hat AI) to alert subagents, a customer-profile agent, line-information agent, router-information agent, wider-configuration agent, and network agent. These agents query the knowledge graph for topology context, pull recent change logs, gather the customer profile and router configuration information, and run RCA using an open source, fine-tuned model served by Red Hat Inference Server.
- **Decision:** The AI leader orchestrator agent identifies the root cause—an outdated router firmware version and misconfigured domain name service (DNS). The agent generates a remediation recommendation and a confidence score. An instant notification is sent to a level-3 engineer who schedules the execution of a Red Hat Ansible Automation Platform playbook.
- **Execution:** The level-3 engineer uses Event-Driven Ansible to execute the remediation playbook. The playbook updates the router firmware, uploads a new router DNS configuration, verifies the customer experience (download/upload speed improvement), and generates an incident report.

The elevated latency use case affects several layers of the architecture: infrastructure, AI, orchestration, agentic reasoning, event transport, and compliance status reporting to the service layer.

The next move: Implementing an intelligent network strategy

The jump to higher levels of network autonomy requires new operating models as well as organizational change, with people and technology working together to achieve AI-augmented operations. Automating high-volume and low-risk tasks is a useful proving ground to validate the value of the AADE cycle. Telecommunication service providers can benefit from mandating the use of open source in all layers to ensure a standardized, interoperable, and transparent approach that capitalizes on the latest innovations from various upstream community projects.

Red Hat's pragmatic framework can take intent and orchestrate AI-driven network optimization at scale within production networks. Its open source foundation gives telecommunication service providers the confidence they need to build an autonomous, intelligent network that delivers business-critical services and applications to meet growing demand.

To learn more about how Red Hat supports telecommunication service providers, contact us today.

North America
1 888 REDHAT1
www.redhat.com

Europe, Middle East, and Africa
00800 7334 2835
europe@redhat.com

Asia Pacific
+65 6490 4200
apac@redhat.com

Latin America
+54 11 4329 7300
info-latam@redhat.com

facebook.com/redhat
x.com/RedHat
linkedin.com/company/red-hat

redhat.com
#3857603_0426